



Sun Cluster 3.1 Release Notes

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 816-5317-10
May 2003, Revision A

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPOUDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



030417@5533



Contents

Sun Cluster 3.1 Release Notes	7
What's New in Sun Cluster 3.1	7
New Features and Functionality	7
Supported Products	9
Restrictions	10
Known Issues and Bugs	15
Incorrect Largefile Status (4419214)	15
Global VxFS File System Lists Block Allocations Differently Than Local VxFS (4449437)	15
Nodes Unable to Bring up qfe Paths (4526883)	15
File Blocks Not Updated Following Writes to Sparse File Holes (4607142)	16
Concurrent use of <code>forcedirectio</code> and <code>mmap(2)</code> may Cause Panics (4629536)	16
Unmounting of a Cluster File System Fails (4656624)	16
Rebooting Puts Cluster Nodes in a Non-Working State (4664510)	17
Dissociating a Plex from a Disk Group Causes Panic (4657088)	17
<code>scvxinstall -i</code> Fails to Install a License Key (4706175)	18
Sun Cluster HA-Siebel Fails to Monitor Siebel Components (4722288)	18
The <code>remove</code> Script Fails to Unregister <code>SUNW.gds</code> Resource Type (4727699)	18
Create IPMP Group Option Overwrites <code>hostname.int</code> (4731768)	18
Using the Solaris <code>shutdown</code> Command May Result in Node Panic (4745648)	19
Administrative Command to Add a Quorum Device to the Cluster Fails (4746088)	19
Path Timeouts When Using <code>ce</code> Adapters on the Private Interconnect (4746175)	19

Siebel Gateway Probe May Time Out When a Public Network Fails (4764204)	20
Flushing Gateway Routes Breaks Per-Node Logical IP Communication (4766076)	20
Unsuccessful Failover Results in Error (4766781)	20
Enabling TCP-Selective Acknowledgments may Cause Data Corruption (4775631)	21
scinstall Incorrectly Shows Some Data Services as Unsupported (4776411)	21
scdidadm Exits With an Error if /dev/rmt is Missing (4783135)	21
Data Corruption When Node Failure Causes the Cluster File System Primary to Die (4804964)	21
Node Hangs After Rebooting When Switchover is in Progress (4806621)	22
File System Panics When Cluster File System is Full (4808748)	22
Cluster Node Hangs While Booting Up (4809076)	22
Removing a Quorum Device Using scconf -rq Causes Cluster Panic (4811232)	23
Mirrored Volume Fails When Using o_EXCL Flag (4820273)	23
Cluster Hangs After a Node is Rebooted During Switchover (4823195)	23
Untranslated Text in the French Locale (4840085)	23
Patches and Required Firmware Levels	24
PatchPro	24
SunSolve Online	24
End-of-Feature-Support Statements	25
Public Network Management (PNM)	25
HASStorage	25
Sun Cluster 3.1 Software Localization	28
Cluster Control Panel (CCP)	29
Installation Tools	30
SunPlex Manager	30
Sun Cluster Module for Sun Management Center	31
Sun Cluster Software	32
Sun Cluster Data Services	32
Sun Cluster 3.1 Documentation	32
Documentation Issues	33
Software Installation Guide	34
SunPlex Manager Online Help	34
System Administration Guide	35
Data Services Collection	35

Man Pages 35

A Sun Cluster Installation and Configuration Worksheets	39
Installation and Configuration Worksheets	40
Local File System Layout Worksheet	42
Cluster and Node Names Worksheet	44
Cluster Interconnect Worksheet	46
Public Networks Worksheet	48
Local Devices Worksheets	50
Disk Device Group Configurations Worksheet	52
Volume Manager Configurations Worksheet	54
Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)	56

Sun Cluster 3.1 Release Notes

This document provides the following information for Sun™ Cluster 3.1 software.

- “What’s New in Sun Cluster 3.1” on page 7
- “Known Issues and Bugs” on page 15
- “Patches and Required Firmware Levels” on page 24
- “End-of-Feature-Support Statements” on page 25
- “Sun Cluster 3.1 Software Localization” on page 28
- “Sun Cluster 3.1 Documentation” on page 32
- “Documentation Issues” on page 33

Note – For information about Sun Cluster 3.1 data services, refer to the *Sun Cluster 3.1 Data Service 5/03 Release Notes*.

What’s New in Sun Cluster 3.1

This section provides information related to new features, functionality, and supported products in Sun Cluster 3.1.

New Features and Functionality

Sun Cluster Security Hardening

Sun Cluster Security Hardening uses the Solaris Operating Environment hardening techniques recommended by the Sun BluePrints™ program to achieve basic security hardening for clusters. The Solaris Security Toolkit automates the implementation of Sun Cluster Security Hardening.

The Sun Cluster Security Hardening documentation is available at <http://www.sun.com/blueprints/0203/817-1079.pdf>. You can also access the article from <http://www.sun.com/software/security/blueprints>. From this URL, scroll down to the Architecture heading to locate the article “Securing the Sun Cluster 3.x Software.” The documentation describes how to secure Sun Cluster 3.1 deployments in a Solaris 8 and Solaris 9 environment. The description includes the use of the Solaris Security Toolkit and other best-practice security techniques recommended by Sun security experts.

Flexible Topologies

Sun Cluster 3.1 software now supports open topologies. You are no longer limited to the storage topologies listed in the *Sun Cluster 3.1 Concepts* document.

Use the following guidelines to configure your cluster.

- Sun Cluster supports a maximum of eight nodes in a cluster, regardless of the storage configurations that you implement.
- A shared storage device can connect to as many nodes as the storage device supports.
- Shared storage devices do not need to connect to all nodes of the cluster. However, these storage devices must connect to at least two nodes.

Diskless Clusters

Sun Cluster 3.1 now supports greater than three-node cluster configurations without shared storage devices. Two-node clusters are still required to have a shared storage device to maintain quorum. This storage device does not need to perform any other function.

Support for Data Service Project Configuration

Data services may now be configured to launch under a Solaris project name when brought online using the RGM—See “Data Service Project Configuration” section in “Key Concepts – Administration and Application Development” in *Sun Cluster 3.1 Concepts Guide* for detailed information about planning project configuration for your data service.

Support for the Solaris Implementation of Internet Protocol (IP) Network Multipathing on Public Networks

For more information on the support for the Solaris implementation of IP network multipathing on public networks, see “Planning the Sun Cluster Configuration” in *Sun Cluster 3.1 Software Installation Guide* and “Administering the Public Network” in *Sun Cluster 3.1 System Administration Guide*.

Set Secondary Nodes for a Disk Device Group

For more information on how to set a desired number of secondary nodes for a disk device group, see “Administering Disk Device Groups” in *Sun Cluster 3.1 System Administration Guide* (refer to the procedures for Setting the Desired Number of Secondaries and Changing Disk Device Group Properties). Additional information can also be found in “Cluster Administration and Application Development” in *Sun Cluster 3.1 Concepts Guide* (See the section on Multi-Ported Disk Failover).

Data Services

For information on data services enhancements, see “What’s New in Sun Cluster 3.1 Data Services 5/03” in *Sun Cluster 3.1 Data Service 5/03 Release Notes*.

Supported Products

This section describes the supported software and memory requirements for Sun Cluster 3.1 software.

- **Operating environment and patches** – Supported Solaris versions and patches are available at the following URL:
<http://sunsolve.sun.com>
For more details, see “Patches and Required Firmware Levels” on page 24.
- **Volume managers**
 - **On Solaris 8** – Solstice DiskSuite™ 4.2.1 and VERITAS Volume Manager 3.2 and 3.5.
 - **On Solaris 9** – Solaris Volume Manager and VERITAS Volume Manager 3.5.

Note – If you are upgrading from VERITAS Volume Manager (VxVM) 3.2 to 3.5, the Cluster Volume Manger (CVM) feature will not be available until you install the CVM license key for version 3.5. In VxVM 3.5, the CVM license key for version 3.2 does not enable CVM and must be upgraded to the CVM license key for version 3.5.

- **File systems** –
 - **On Solaris 8** – Solaris UFS and VERITAS File System 3.4 and 3.5.
 - **On Solaris 9** – Solaris UFS and VERITAS File System 3.5.
- **Data services (agents)** – For information on supported data services, see *Sun Cluster 3.1 Data Service 5/03 Release Notes*.

Note – Sun Cluster 3.0 data services can run on Sun Cluster 3.1, except as noted in “Running Sun Cluster HA for Oracle 3.0 on Sun Cluster 3.1” on page 14.

- **Memory Requirements** – Sun Cluster 3.1 software requires extra memory beyond what is configured for a node under a normal workload. The extra memory equals 128 Mbytes plus ten percent. For example, if a standalone node normally requires 1 Gbyte of memory, you need an extra 256 Mbytes to meet memory requirements.
- **RSMAPI** –Sun Cluster 3.1 software supports the Remote Shared Memory Application Programming Interface (RSMAPI) on RSM-capable interconnects, such as PCI-SCI.

Restrictions

The following restrictions apply to the Sun Cluster 3.1 release:

- **svc_default_stksize and lwp_default_stksize parameters** –Set the `rpcmod:svc_default_stksize` parameter to `0x8000` and the `lwp_default_stksize` parameter to `0x6000` in the `/etc/system` file, to avoid stack overflow.

Note – If any VxFS package or patch is added, make sure that the settings for these parameters in the `/etc/system` file match the values shown above.

- **local-mac-address? variable** – The `local-mac-address?` variable must have a value of `true` for Ethernet adapters. This is a reversal of the Sun Cluster 3.0 software requirement, which was to set this variable to a value of `false`.
- **Remote Shared Memory (RSM) transport types** – These transport types are mentioned in the documentation, but are not supported. If you use the RSMAPI, specify `dlpi` as the transport type.
- **Scalable Coherent Interface (SCI)** – The SBus SCI interface is not supported as a cluster interconnect. However, the PCI-SCI interface is supported.
- **Logical network interfaces** – These interfaces are reserved for use by Sun Cluster 3.1 software.
- **Disk path monitoring** – Only active disk paths are monitored on the current primary node for failures by Sun Cluster software. You must monitor disk paths manually to avoid double failures or loss of path to a quorum device.
- **Storage devices with more than two physical paths to the enclosure** – More than two paths are not supported except on the following: Sun StorEdge™ A3500, for which two paths are supported to each of two nodes, any device that supports Sun StorEdge Traffic Manager, and EMC storage devices that use EMC PowerPath

software.

- **SunVTS™** – Not supported.
- **Multihost tape, CD-ROM, and DVD-ROM** – Not supported.
- **Loopback File System** – Sun Cluster 3.1 software does not support the use of the loopback file system (LOFS) on cluster nodes.
- **Running client applications on the cluster nodes** – Client applications that run on cluster nodes should not map to logical IP addresses of an HA data service. During failover, these logical IP addresses might go away, leaving the client without a connection.
- **Running high-priority process scheduling classes on cluster nodes** – Not supported. Processes that run in the time-sharing scheduling class with a high priority, or processes that run in the real-time scheduling class should not be run on cluster nodes. Sun Cluster 3.1 software relies on kernel threads that do not run in the real-time scheduling class. Other time-sharing processes that run at higher-than-normal priority or real-time processes can prevent the Sun Cluster kernel threads from acquiring needed CPU cycles.
- **Upgrade from Solaris 8 to Solaris 9** - Upgrade from Solaris 8 to Solaris 9 software on a Sun Cluster configuration is not supported. You can only upgrade to subsequent, compatible versions of the Solaris 8 operating environment. To run Sun Cluster 3.1 software on the Solaris 9 operating environment, you must perform a new installation of the Solaris 9 version of Sun Cluster 3.1 software after the nodes are installed with Solaris 9 software.
- **IPv6** - Not supported.
- **SNDR cannot be used with HAStoragePlus** - Currently, SNDR can only be used with HAStorage. This restriction only applies to the light weight resource group that includes the logical host SNDR is using for replication. Application resource groups can still use HAStoragePlus with SNDR. You can use failover filesystem with HAStoragePlus and SNDR by using HAStorage for the SNDR resource group, and HAStoragePlus for the application resource group, where the HAStorage and HAStoragePlus resources point at the same underlying DCS device. A patch is being developed to enable SNDR to work with HAStoragePlus.
- **Mounting options** - (1) You cannot remount a file system with the `directio` mount option added at remount time, and (2) you cannot set the `directio` mount option on a single file by using the `directio ioctl`.
- **License key** - The license key can only be installed with the interactive form or with the `scvxinstall -e` option.
- **Other restrictions** - For other known problems or restrictions, see “Known Issues and Bugs” on page 15.

Service and Application Restrictions

- Sun Cluster 3.1 software can only provide service for those data services that are either supplied with the Sun Cluster product or set up with the Sun Cluster data services API.
- Sun Cluster software currently does not have an HA Data Service for the `sendmail(1M)` subsystem. It is permitted to run `sendmail` on the individual cluster nodes, but the `sendmail` functionality will not be highly available, including the functionality of mail delivery and mail routing, queuing, or retry.
- Do not configure cluster nodes as routers (gateways). If the system goes down, the clients cannot find an alternate router and cannot recover.
- Do not configure cluster nodes as NIS or NIS+ servers. However, cluster nodes can be NIS or NIS+ clients.
- Do not use a Sun Cluster configuration to provide a highly available boot or install service on client systems.
- Do not use a Sun Cluster 3.1 configuration to provide an `rarpd` service.

Hardware Restrictions

- Alternate Pathing (AP) is not supported.
- If you are using a Sun Enterprise™ 420R server with a PCI card in slot J4701, the motherboard must be at dash-level 15 or higher (501-5168-15 or higher). To find the motherboard part number and revision level, look at the edge of the board closest to PCI slot 1.
- System panics have been observed in clusters when UDWIS I/O cards are used in slot 0 of a board in a Sun Enterprise 10000 server; do not install UDWIS I/O cards in slot 0 of a board in this server.

Volume Manager Restrictions

- If you are upgrading from VERITAS Volume Manager (VxVM) 3.2 to 3.5, the Cluster Volume Manger (CVM) feature will not be available until you install the CVM license key for version 3.5. In VxVM 3.5, the CVM license key for version 3.2 does not enable CVM and must be upgraded to the CVM license key for version 3.5.
- In Solstice DiskSuite/Solaris Volume Manager configurations that use mediators, the number of mediator hosts configured for a diskset must be exactly two.
- DiskSuite Tool (Solstice DiskSuite `metatool`) and the Enhanced Storage module of Solaris Management Console (Solaris Volume Manager) are not compatible with Sun Cluster 3.1 software.
- Use of VxVM Dynamic Multipathing (DMP) with Sun Cluster 3.1 software to manage multiple paths from the same node is not supported. From VxVM 3.2 onward it is no longer possible to disable the installation of DMP. But having it in

the I/O stack on systems with only a single path per node poses no problems. However, if you use VxVM in a configuration with multiple paths per node, then you must use another multipathing solution, such as MPxIO or EMC PowerPath.

- Simple root disk groups (`rootdg` created on a single slice of the root disk) are not supported as disk types with VxVM on Sun Cluster 3.1 software.
- Software RAID 5 is not supported.

Cluster File System Restrictions

- Quotas are not supported by Sun Cluster file systems.
- The command `umount -f` behaves in the same manner as the `umount` command without the `-f` option. It does not support forced unmounts.
- The command `unlink (1M)` is not supported on non-empty directories.
- The command `lockfs -d` is not supported. Use `lockfs -n` as a workaround.
- The cluster file system does not support any of the file-system features of Solaris software by which one would put a communication end-point in the file-system name space. Therefore, although you can create a UNIX domain socket whose name is a path name into the cluster file system, the socket would not survive a node failover. In addition, any fifos or named pipes you create on a cluster file system would not be globally accessible, nor should you attempt to use `fattach` from any node other than the local node.
- It is not supported to execute binaries off file systems mounted by using the `forcedirectio` mount option.

VxFS Restrictions

- The following VxFS features are not supported in a Sun Cluster 3.1 configuration.
 - Quick I/O
 - Snapshots
 - Storage checkpoints
 - Cache advisories (these can be used, but the effect will be observed on the given node only)
 - VERITAS CFS (requires VERITAS cluster feature and VCS)

All other VxFS features and options that are supported in a cluster configuration are supported by Sun Cluster 3.1 software. See VxFS documentation and man pages for details about VxFS options that are or are not supported in a cluster configuration.

- The following VxFS-specific mount options are not supported in a Sun Cluster 3.1 configuration.
 - `convosync` (Convert `O_SYNC`)

- mincache
- qlog, delaylog, tmplog
- For information about administering VxFS cluster file systems in a Sun Cluster configuration, see “Administering Cluster File Systems Overview” in *Sun Cluster 3.1 System Administration Guide*.

Internet Protocol (IP) Network Multipathing Restrictions

This section identifies any restrictions on using IP Network Multipathing that apply only in a Sun Cluster 3.1 environment, or are different than information provided in the Solaris documentation for IP Network Multipathing.

- IPv6 is not supported.
- All public network adapters must be in IP Network Multipathing groups.
- In `/etc/default/mpathd` do not change `TRACK_INTERFACES_ONLY_WITH_GROUPS` from yes to no.
- For knownbugs and issues, see “Create IPMP Group Option Overwrites `hostname.int` (4731768)” on page 18 .

Most procedures, guidelines, and restrictions identified in the Solaris documentation for IP Network Multipathing are the same in a cluster or non-cluster environment. Therefore, see the appropriate Solaris document for additional information about IP Network Multipathing restrictions.

Operating Environment Release	For Instructions, Go To...
Solaris 8 operating environment	<i>IP Network Multipathing Administration Guide</i>
Solaris 9 operating environment	“IP Network Multipathing Topics” in <i>System Administration Guide: IP Series</i>

Data Service Restrictions

There are no restrictions that apply to all data services. For information about restrictions for specific data services, see *Sun Cluster 3.1 Data Service 5/03 Release Notes*.

Running Sun Cluster HA for Oracle 3.0 on Sun Cluster 3.1

The Sun Cluster HA for Oracle 3.0 data service can run on Sun Cluster 3.1 only when used with the following versions of the Solaris operating environment:

- Solaris 8, 32-bit version
- Solaris 8, 64-bit version
- Solaris 9, 32-bit version

Note – The Sun Cluster HA for Oracle 3.0 data service *cannot* run on Sun Cluster 3.1 when used with the 64-bit version of Solaris 9.

Known Issues and Bugs

The following known issues and bugs affect the operation of the Sun Cluster 3.1 release. For the most current information, see the online *Sun Cluster 3.1 Release Notes Supplement* at <http://docs.sun.com>.

Incorrect Largefile Status (4419214)

Problem Summary: The `/etc/mnttab` file does not show the most current largefile status of a globally mounted VxFS filesystem.

Workaround: Use the `fsadm` command to verify the filesystem largefile status, instead of the `/etc/mnttab` entry.

Global VxFS File System Lists Block Allocations Differently Than Local VxFS (4449437)

Problem Summary: For a given file size, global VxFS file system appears to allocate more disk blocks than the local VxFS file system.

Workaround: Unmounting and mounting the filesystem eliminates the extra disk blocks that were reported as allocated to the given file.

Nodes Unable to Bring up qfe Paths (4526883)

Problem Summary: Sometimes, private interconnect transport paths ending at a qfe adapter fail to come online.

Workaround: Follow the steps shown below:

1. Using `scstat -w`, identify the adapter that is at fault. The output will show all transport paths with that adapter as one of the path endpoints in the `faulted` or the `waiting` states.

2. Use `scsetup` to remove from the cluster configuration all the cables connected to that adapter.
3. Use `scsetup` again to remove that adapter from the cluster configuration.
4. Add back the adapter and the cables.
5. Verify if the paths appear. If the problem persists, repeat steps 1–5 a few times.
6. Verify if the paths appear. If the problem still persists, reboot the node with the at-fault adapter. Before the node is rebooted, make sure that the remaining cluster has enough quorum votes to survive the node reboot.

File Blocks Not Updated Following Writes to Sparse File Holes (4607142)

Problem Summary: A file's block count is not always consistent across cluster nodes following block-allocating write operations within a sparse file. For a cluster file system layered on UFS (or VxFS 3.4), the block inconsistency across cluster nodes disappears within 30 seconds or so.

Workaround: File metadata operations which update the `inode` (`touch`, etc.) should synchronize the `st_blocks` value so that subsequent metadata operations will ensure consistent `st_blocks` values.

Concurrent use of `forcedirectio` and `mmap(2)` may Cause Panics (4629536)

Problem Summary: Using the `forcedirectio` mount option and the `mmap(2)` function concurrently might cause data corruption, system hangs, or panics.

Workaround: Observe the following restrictions:

- Do not remount a file system with the `directio` mount option added at remount time.
- Do not set the `directio` mount option on a single file by using the `directio` `ioctl`.

If there is a need to use `directio`, mount the whole file system with `directio` options.

Unmounting of a Cluster File System Fails (4656624)

Problem Summary: The unmounting of a cluster file system fails sometimes even though the `fuser` command shows that there are no users on any node.

Workaround: Retry the unmounting after all asynchronous I/O to the underlying file system has been completed.

Rebooting Puts Cluster Nodes in a Non-Working State (4664510)

Problem Summary: After powering off one of the Sun StorEdge T3 Arrays and running `scshut down`, rebooting both nodes puts the cluster in a non-working state.

Workaround: If half the replicas are lost, perform the following steps:

1. Ensure the cluster is in cluster mode.
2. Forcibly import the diskset.

```
# metaset -s set-name -f -C take
```

3. Delete the broken replicas.

```
# metadb -s set-name -fd /dev/did/dsk/dNsX
```

4. Release the diskset.

```
# metaset -s set-name -C release
```

The file system can now be mounted and used. However, the redundancy in the replicas has not been restored. If the other half of replicas is lost, then there will be no way to restore the mirror to a sane state.

5. Recreate the databases after the above repair procedure is applied.

Dissociating a Plex from a Disk Group Causes Panic (4657088)

Problem Summary: Dissociating or detaching a plex from a disk group under Sun Cluster may panic the cluster node with following panic string:

```
panic[cpu2]/thread=30002901460: BAD TRAP: type=31 rp=2a101b1d200  
addr=40 mmu_fsr=0 occurred in module "vxfs" due to a NULL pointer  
dereference
```

Workaround: Before dissociating or detaching a plex from a disk group, unmount the corresponding file system.

scvxinstall -i Fails to Install a License Key (4706175)

Problem Summary: The `scvxinstall -i` command accepts a license key with the `-L` option. However, the key is ignored and does not get installed.

Workaround: Do not provide a license key with the `-i` form of `scvxinstall`. The key will not be installed. The license keys should be installed with the interactive form or with the `-e` option. Before proceeding with the encapsulation of root, examine the license requirements and provide the desired keys either with the `-e` option or in the interactive form.

Sun Cluster HA-Siebel Fails to Monitor Siebel Components (4722288)

Problem Summary: The Sun Cluster HA-Siebel agent will not monitor individual Siebel components. If failure of a Siebel component is detected, only a warning message would be logged in `syslog`.

Workaround: Restart the Siebel server resource group in which components are offline using the command `scswitch -R -h node-g resource_group`.

The remove Script Fails to Unregister SUNW.gds Resource Type (4727699)

Problem Summary: The `remove` script fails to unregister `SUNW.gds` resource type and displays the following message:

```
Resource type has been un-registered already.
```

Workaround: After using the `remove` script, manually unregister `SUNW.gds`. Alternatively, use the `scsetup` command or the SunPlex Manager.

Create IPMP Group Option Overwrites hostname.int (4731768)

Problem Summary: The `Create IPMP group` option in SunPlex Manager should only be used with adapters that are not already configured. If an adapter is already configured with an IP address, the adapter must be manually configured for IPMP.

Workaround: The `Create IPMP group` option in SunPlex Manager must be used only with adapters that are not already configured. If an adapter is already configured with an IP address, the adapter should be manually configured using Solaris IPMP management tools.

Using the Solaris shutdown Command May Result in Node Panic (4745648)

Problem Summary: Using the Solaris `shutdown` command or similar commands (for example, `uadmin`) to bring down a cluster node may result in node panic and display the following message:

```
CMM: Shutdown timer expired. Halting.
```

Workaround: Contact your Sun service representative for support. The panic is necessary to provide a guaranteed safe way for another node in the cluster to take over the services that were being hosted by the shutting-down node.

Administrative Command to Add a Quorum Device to the Cluster Fails (4746088)

Problem Summary: If a cluster has the minimum votes required for quorum, an administrative command to add a quorum device to the cluster fails with the following message:

```
Cluster could lose quorum.
```

Workaround: Contact your Sun service representative for support.

Path Timeouts When Using ce Adapters on the Private Interconnect (4746175)

Problem Summary: Clusters using `ce` adapters on the private interconnect may notice path timeouts and subsequent node panics if one or more cluster nodes have more than four processors.

Workaround: Set the `ce_taskq_disable` parameter in the `ce` driver by adding `set ce:ce_taskq_disable=1` to `/etc/system` file on all cluster nodes and then rebooting the cluster nodes. This ensures that heartbeats (and other packets) are always delivered in the interrupt context, eliminating path timeouts and the subsequent node panics. Quorum considerations should be observed while rebooting cluster nodes.

Siebel Gateway Probe May Time Out When a Public Network Fails (4764204)

Problem Summary: Failure of a public network may cause the Siebel gateway probe to time out and eventually cause the Siebel gateway resource to go offline. This may occur if the node on which the Siebel gateway is running has a path beginning with `/home` which depends on network resources such as NFS and NIS. Without the public network, the Siebel gateway probe hangs while trying to open a file on `/home`, causing the probe to time out.

Workaround: Complete the following steps for all nodes of the cluster which can host the Siebel gateway.

1. Ensure that the `passwd`, `group`, and `project` entries in `/etc/nsswitch.conf` refer only to *files* and not to *nis*.
2. Ensure that there are no NFS or NIS dependencies for any path starting with `/home`. You may have either a locally mounted `/home` path or rename the `/home` mount point to `/export/home` or some name which does not start with `/home`.
3. In the `/etc/auto_master` file, comment out the line containing the entry `+auto_master`. Also comment out any `/home` entries using `auto_home`.
4. In `etc/auto_home`, comment out the line containing `+auto_home`.

Flushing Gateway Routes Breaks Per-Node Logical IP Communication (4766076)

Problem Summary: To provide highly available, per-node, logical IP communication over a private interconnect, Sun Cluster software relies on gateway routes on the cluster nodes. Flushing the gateway routes will break the per-node logical IP communication.

Workaround: Reboot the cluster nodes where the routes were inadvertently flushed. To restore the gateway routes, it is sufficient to reboot the cluster nodes one at a time. Per-node logical IP communication will remain broken until the routes have been restored. Quorum considerations must be observed while rebooting cluster nodes.

Unsuccessful Failover Results in Error (4766781)

Problem Summary: An unsuccessful failover/switchover of a file system might leave the file system in an errored state.

Workaround: Unmount and remount the file system.

Enabling TCP-Selective Acknowledgments may Cause Data Corruption (4775631)

Problem Summary: Enabling TCP-selective acknowledgements on cluster nodes may cause data corruption.

Workaround: No user action is required. To avoid causing data corruption on the global file system, do not reenable TCP selective acknowledgements on cluster nodes.

`scinstall` Incorrectly Shows Some Data Services as Unsupported (4776411)

Problem Summary: `scinstall` incorrectly shows that the following data services are not supported on Solaris 9:

- Sun Cluster HA for SAP
- Sun Cluster HA for SAP liveCache

Workaround: Solaris 8 and 9 support both Sun Cluster HA for SAP and Sun Cluster HA for SAP liveCache; ignore the unsupported feature list in `scinstall`.

`scdidadm` Exits With an Error if `/dev/rmt` is Missing (4783135)

Problem Summary: The current implementation of `scdidadm(1M)` relies on the existence of both `/dev/rmt` and `/dev/(r) dsk` to successfully execute `scdiadm -r`. Solaris installs both, regardless of the existence of the actual underlying storage devices. If `/dev/rmt` is missing, `scdidadm` exits with the following error:

```
Cannot walk /dev/rmt" during execution of 'scdidadm -r.
```

Workaround: On any node where `/dev/rmt` is missing, use `mkdir` to create a directory `/dev/rmt`. Then, run `scgdevs` from one node.

Data Corruption When Node Failure Causes the Cluster File System Primary to Die (4804964)

Problem Summary: Data corruption may occur with Sun Cluster 3.x systems running patches 113454-04, 113073-02 and 113276-02 (or a subset of these patches). The problem only occurs with globally mounted UFS file systems. The data corruption results in missing data (that is, you will see zero's where data should exist), and the amount of missing data is always a multiple of a disk block. The data loss can occur any time a

node failure causes the cluster file system primary to die soon after the cluster file system client completes— or reports that it has just completed—a write operation. The period of vulnerability is limited and does not occur every time.

Workaround: Use the `-o syncdir` mount option to force UFS to use synchronous UFS log transactions.

Node Hangs After Rebooting When Switchover is in Progress (4806621)

Problem Summary: If a device group switchover is in progress when a node joins the cluster, the joining node and the switchover operation may hang. Any attempts to access any device service will also hang. This is more likely to happen on a cluster with more than two nodes and if the file system mounted on the device is a VxFS file system.

Workaround: To avoid this situation, do not initiate device group switchovers while a node is joining the cluster. If this situation occurs, then all the cluster nodes must be rebooted to restore access to device groups.

File System Panics When Cluster File System is Full (4808748)

Problem Summary: When a cluster file system is full, there are instances where the filesystem might panic with one of the following messages: 1)

```
assertion failed: cur_data_token & PXFS_WRITE_TOKEN or  
PXFS_READ_TOKEN or 2)
```

```
vp->v_pages == NULL. These panics are intended to prevent data  
corruption when a filesystem is full.
```

Workaround: To reduce the likelihood of this problem, use a cluster file system with UFS as far as possible. It is extremely rare for one of these panics to occur when using a cluster file system with UFS, but the risk is greater when using a cluster file system with VxFS.

Cluster Node Hangs While Booting Up (4809076)

Problem Summary: When a device service switchover request, using `scswitch -z -D <device-group> -h <node>`, is concurrent with a node reboot and there are global file systems configured on the device service, the global file systems might become unavailable and subsequent configuration changes involving any device service or global file system may also hang. Additionally, subsequent cluster node joins might hang.

Workaround: Recovery requires a reboot of all the cluster nodes.

Removing a Quorum Device Using `scconf -rq` Causes Cluster Panic (4811232)

Problem Summary: If you execute the `scconf -rq` command to remove a quorum device in a vulnerable configuration, all nodes of the cluster will panic with the message

```
CMM lost operational quorum.
```

Workaround: To remove a quorum device from a cluster, first check the output of `scstat -q`. If the quorum device is listed as having more than one vote in the Present column, then the device should first be put into maintenance mode using `scconf -cq globaldev=QD, maintstate`. After the command completes and the quorum device is shown in `scstat -q` as having 0 votes present, the device can be removed using `scconf -rq`.

Mirrored Volume Fails When Using `O_EXCL` Flag (4820273)

Problem Summary: If Solstice DiskSuite/Solaris Volume Manager is being used and a mirrored volume is opened with `O_EXCL` flag, the failover of the device group containing this volume will fail. This will panic the new device group primary, when the volume is first accessed after the failover.

Workaround: When using Solstice DiskSuite/Solaris Volume Manager, do not open mirrored volumes with `O_EXCL` flag.

Cluster Hangs After a Node is Rebooted During Switchover (4823195)

Problem Summary: If a device service failover request is concurrent with a node reboot or a node join, and there are cluster file systems configured on the device service, the cluster file systems might become unavailable and subsequent configuration changes involving any device service or cluster file system may also hang. Additionally, subsequent cluster node joins might hang.

Workaround: Recovery requires a reboot of all the cluster nodes.

Untranslated Text in the French Locale (4840085)

Problem Summary: Some untranslated text appears when using the SunPlex Manager to install Sun Cluster in the French locale.

Workaround: This error does not affect SunPlex Manager's functionality. You may either ignore the untranslated text or set your browser's language to English to avoid mixed translation.

Patches and Required Firmware Levels

This section provides information about patches for Sun Cluster configurations.

Note – You must be a registered SunSolve™ user to view and download the required patches for the Sun Cluster product. If you do not have a SunSolve account, contact your Sun service representative or sales engineer, or register online at <http://sunsolve.sun.com>.

PatchPro

PatchPro is a patch-management tool designed to ease the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides a Sun Cluster-specific Interactive Mode tool to make the installation of patches easier and an Expert Mode tool to maintain your configuration with the latest set of patches. Expert Mode is especially useful for those who want to get all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to <http://www.sun.com/PatchPro/>, click on "Sun Cluster," then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

SunSolve Online

The SunSolve™ Online Web site provides 24-hour access to the most up-to-date information regarding patches, software, and firmware for Sun products. Access the SunSolve Online site at <http://sunsolve.sun.com> for the most current matrixes of supported software, firmware, and patch revisions.

You can find Sun Cluster 3.1 patch information by using the Info Docs. To view the Info Docs, log on to SunSolve and access the Simple search selection from the top of the main page. From the Simple Search page, click on the Info Docs box and type **Sun Cluster 3.1** in the search criteria box. This will bring up the Info Docs page for Sun Cluster 3.1 software.

Before you install Sun Cluster 3.1 software and apply patches to a cluster component (Solaris operating environment, Sun Cluster software, volume manager or data services software, or disk hardware), review the Info Docs and any README files that accompany the patches. All cluster nodes must have the same patch level for proper cluster operation.

For specific patch procedures and tips on administering patches, see the *Sun Cluster 3.1 System Administration Guide*.

End-of-Feature-Support Statements

Public Network Management (PNM)

Public Network Management (PNM) is not supported in Sun Cluster 3.1. Software Network adapter monitoring and failover for Sun Cluster software is instead performed by the Solaris implementation of Internet Protocol (IP) Network Multipathing. See “What’s New in Sun Cluster 3.1” on page 7.

HAStorage

HAStorage might not be supported in a future release of Sun Cluster software. Near-equivalent functionality is supported by HAStoragePlus. Complete one of the following procedures to migrate from HAStorage to HAStoragePlus.

▼ How to Upgrade from HAStorage to HAStoragePlus When Using Device Groups or Cluster File Systems

HAStorage might not be supported in a future release of Sun Cluster software. Equivalent functionality is supported by HAStoragePlus. To upgrade from HAStorage to HAStoragePlus when you use cluster file systems or device groups, complete the following steps.

The following example uses a simple HA-NFS resource active with HAStorage. The ServicePaths are the disk group `nfsdg` and the `AffinityOn` property is `TRUE`. Furthermore, the HA-NFS Service has `Resource_Dependencies` set to the HAStorage resource.

1. **Remove the dependencies the application resources has on HAStorage resource.**

```
# scrgadm -c -j nfsserver-rs -y Resource_Dependencies=""
```

2. Disable the HAStorage resource.

```
# scswitch -n -j nfs1storage-rs
```

3. Remove the HAStorage resource from the application resource group.

```
# scrgadm -r -j nfs1storage-rs
```

4. Unregister the HAStorage resource type.

```
# scrgadm -r -t SUNW.HAStorage
```

5. Register the HAStoragePlus resource type.

```
# scrgadm -a -t SUNW.HAStoragePlus
```

6. Create the HAStoragePlus resource.

To specify a file-system mount point, input the following text.

```
# scrgadm -a -j nfs1-hastp-rs -g nfs1-rg -t \  
SUNW.HAStoragePlus -x FilesystemMountPoints=/global/nfsdata -x \  
AffinityOn=True
```

To specify global device paths, input the following text.

```
# scrgadm -a -j nfs1-hastp-rs -g nfs1-rg -t \  
SUNW.HAStoragePlus -x GlobalDevicePaths=nfsdg -x AffinityOn=True
```

Note – Instead of using the `ServicePaths` property for HAStorage, you must use the `GlobalDevicePaths` or `FilesystemMountPoints` property for HAStoragePlus. The `FilesystemMountPoints` extension property must match the sequence specified in the `/etc/vfstab` file.

7. Enable the HAStoragePlus resource.

```
# scswitch -e -j nfs1-hastp-rs
```

8. Set up the dependencies between the application server and HAStoragePlus.

```
# scrgadm -c -j nfsserver-rs -y \  
Resource_Dependencies=nfs1=hastp-rs
```

How to Upgrade from HAStorage with Cluster File Systems to HAStoragePlus with Failover Filesystem

HAStorage might not be supported in a future release of Sun Cluster. Equivalent functionality is supported by HAStoragePlus. To upgrade from HAStorage to HAStoragePlus when using Failover Filesystem (FFS), complete the following steps.

The following example uses a simple NFS service active with HAStorage. The ServicePaths are the diskgroup nfsdg and the AffinityOn property is TRUE. Furthermore, the HA-NFS service has Resource_Dependencies set to the HAStorage resource.

1. Remove the dependencies the application resource has on HAStorage.

```
# scrgadm -c -j nfserver-rs -y Resource_Dependencies=""'
```

2. Disable the HAStorage resource.

```
# scswitch -n -j nfs1storage-rs
```

3. Remove the HAStorage resource from the application resource group.

```
# scrgadm -r -j nfs1storage-rs
```

4. Unregister the HAStorage resource type.

```
# scrgadm -r -t SUNW.HAStorage
```

5. Modify the /etc/vfstab file to remove the global flag and change mount at boot to no. This should be done on all nodes which are potential primaries for the resource group.

6. Register the HAStoragePlus resource type.

```
# scrgadm -a -t SUNW.HAStoragePlus
```

7. Create the HAStoragePlus resource.

To specify a file-system mount point, input the following text.

```
# scrgadm -a -j nfs1-hastp-rs -g nfs1-rg -t \  
SUNW.HAStoragePlus -x FilesystemMountPoints=/global/nfsdata -x \  
AffinityOn=True
```

To specify global device paths, input the following text.

```
# scrgadm -a -j nfs1-hastp-rs -g nfs1-rg -t \  
SUNW.HAStoragePlus -x GlobalDevicePaths=nfsdg -x AffinityOn=True
```

Note – Instead of using the ServicePaths property for HAStorage, you must use the GlobalDevicePaths or FilesystemMountPoints property for HAStoragePlus. The FilesystemMountPoints extension property must match the sequence specified in the /etc/vfstab file.

8. Switch the application resource group offline.

```
# scswitch -F -g nfs1-rg
```

9. Disable the application resource.

```
# scswitch -n -j nfserver-rs
```

10. Unmount the CFS file systems.

11. Enable the HAStoragePlus resource.

```
# scswitch -e -j nfs1-hastp-rs
```

12. Bring the application resource group online on a given host.

```
# scswitch -z -g nfs1-rg -h hostname
```

13. Set up the dependencies between the application resource and HAStoragePlus.

```
# scrgadm -c -j nfserver-rs -y \  
Resource_Dependencies=nfs1=hastp-rs
```

14. Enable the application resource.

```
# scswitch -e -j nfs1-hastp-rs
```

Sun Cluster 3.1 Software Localization

Localization is available for selected Sun Cluster software components in the following languages:

Language	Localized Sun Cluster Component
French	Installation Cluster Control Panel (CCP) Sun Cluster Software Sun Cluster Data Services Sun Cluster module for Sun Management Center SunPlex Manager

Language	Localized Sun Cluster Component
Japanese	Installation Cluster Control Panel (CCP) Sun Cluster Software Sun Cluster Data Services Sun Cluster module for Sun Management Center SunPlex Manager Sun Cluster man pages Cluster Control Panel man pages Sun Cluster Data Service messages man pages
Simplified Chinese	Sun Cluster module for Sun Management Center SunPlex Manager
Traditional Chinese	Sun Cluster module for Sun Management Center (online help only) SunPlex Manager (online help only)
Korean	Sun Cluster module for Sun Management Center (online help only) SunPlex Manager (online help only)

The following sections provide instructions on how to install the localization packages for various Sun Cluster components:

- “Cluster Control Panel (CCP)” on page 29
- “Installation Tools” on page 30
- “SunPlex Manager” on page 30
- “Sun Cluster Module for Sun Management Center” on page 31
- “Sun Cluster Software” on page 32
- “Sun Cluster Data Services” on page 32

Cluster Control Panel (CCP)

To use the localized Cluster Control Panel (CCP), you must first install the following packages on your administrative console by using the `pkgadd(1M)` command.

Language	Package Name	Package Description
French	SUNWFCCON	French Sun Cluster Console
Japanese	SUNWJCCON	Japanese Sun Cluster Console

Language	Package Name	Package Description
Simplified Chinese	SUNWcccon	Simplified Chinese Sun Cluster Console

Installation Tools

To use the localized `scinstall(1M)` utility to install Sun Cluster 3.1 software, install the following packages on the cluster nodes by using the `pkgadd(1M)` command before you run `scinstall`.

Language	Package Name	Package Description
French	SUNWfsc	French Sun Cluster messages
Japanese	SUNWjsc	Japanese Sun Cluster messages
	SUNWjscman	Japanese Sun Cluster man pages

To use the localized SunPlex Manager to install Sun Cluster 3.1 software, see "SunPlex Manager" on page 30 for more information.

SunPlex Manager

To use the localized SunPlex Manager, the following packages are required on the cluster nodes.

Language	Package Name	Package Description
French	SUNWfsc	French Sun Cluster messages
	SUNWfscvw	French SunPlex Manager online help
Japanese	SUNWjsc	Japanese Sun Cluster messages
	SUNWjscvw	Japanese SunPlex Manager online help
Simplified Chinese	SUNWcsc	Simplified Chinese Sun Cluster messages
	SUNWcscvw	Simplified Chinese SunPlex Manager online help
Traditional Chinese	SUNWhscvw	Traditional Chinese SunPlex Manager online help
Korean	SUNWkscvw	Korean SunPlex Manager online help

After you install the localized SunPlex Manager packages, set your browser's language preference. If you are using Netscape, you can check and change browser languages by performing the following steps:

1. Start Netscape.
2. Select Edit > Preferences from the main menu.
3. Select Navigator > Languages from the Preferences dialog box.
4. Click Add, then select the language you want from the Add Language dialog box.
5. Click OK.

Sun Cluster Module for Sun Management Center

To use the localized Sun Cluster module for Sun Management Center, install the following packages to the Sun Management Center server layer by using the `pkgadd(1M)` command.

Language	Package Name	Package Description
French	SUNWfscsv	French Sun Cluster SyMON server add-on
Japanese	SUNWjscsv	Japanese Sun Cluster SyMON server add-on
Simplified Chinese	SUNWcscsv	Simplified Chinese Sun Cluster SyMON server add-on

To use the localized online help on the Sun Cluster module for Sun Management Center, install the following packages to the Sun Management Center console layer by using the `pkgadd(1M)` command.

Language	Package Name	Package Description
French	SUNWfscshl	French Sun Cluster SyMON modules
Japanese	SUNWjscshl	Japanese Sun Cluster SyMON modules
Simplified Chinese	SUNWcscshl	Simplified Chinese Sun Cluster SyMON modules
Traditional Chinese	SUNWhscshl	Traditional Chinese Sun Cluster SyMON modules
Korean	SUNWkscshl	Korean Sun Cluster SyMON modules

Sun Cluster Software

The following Sun Cluster localization packages will be automatically installed on the cluster node when you install or upgrade to Sun Cluster 3.1.

Language	Package Name	Package Description
French	SUNWfsc	French Sun Cluster messages
Japanese	SUNWjsc	Japanese Sun Cluster messages
	SUNWjscman	Japanese Sun Cluster man pages
Simplified Chinese	SUNWcsc	Simplified Chinese Sun Cluster messages

Sun Cluster Data Services

When you install or upgrade to Sun Cluster 3.1, the localization packages will be automatically installed for the data services you have selected. For more information, see *Sun Cluster 3.1 Data Service 5/03 Release Notes*.

Sun Cluster 3.1 Documentation

The complete Sun Cluster 3.1 user documentation set is available in PDF and HTML format on both the Sun Cluster 3.1 CD-ROM and the Sun Cluster 3.1 Agents CD-ROM. AnswerBook2™ server software is not needed to read Sun Cluster 3.1 documentation. See the `index.html` file at the top level of either CD-ROM for more information. This `index.html` file enables you to read the PDF and HTML manuals directly from the disc and to access instructions to install the documentation packages.

Note – The `SUNWsdocs` package must be installed before you install any Sun Cluster documentation packages. You can use `pkgadd` to install the `SUNWsdocs` package. The `SUNWsdocs` package is located in the `SunCluster_3.1/Sol_N/Packages/` directory of the Sun Cluster 3.1 CD-ROM, where *N* is either 8 for Solaris 8 or 9 for Solaris 9. The `SUNWsdocs` package is also automatically installed when you run the installer from the Solaris 9 Documentation CD.

The Sun Cluster 3.1 documentation set consists of the following collections:

- The Sun Cluster 3.1 Software Collection, which includes the following manuals:
Sun Cluster 3.1 Concepts Guide

Sun Cluster 3.1 Data Services Developer's Guide

Sun Cluster 3.1 Error Messages Guide

Sun Cluster 3.1 Software Installation Guide

Sun Cluster 3.1 System Administration Guide

- The Sun Cluster 3.x Hardware Administration Collection, which includes the following manuals:

Sun Cluster 3.x Hardware Administration Manual

Sun Cluster 3.x With Sun StorEdge 3310 Array Manual

Sun Cluster 3.x With Sun StorEdge 3900 or 6900 Series System Manual

Sun Cluster 3.x With Sun StorEdge 9900 Series Storage Device Manual

Sun Cluster 3.x With Sun StorEdge A1000 or Netra st A1000 Array Manual

Sun Cluster 3.x With Sun StorEdge A3500/A3500FC System Manual

Sun Cluster 3.x With Sun StorEdge A5x00 Array Manual

Sun Cluster 3.x With Sun StorEdge D1000 or Netra st D1000 Disk Array Manual

Sun Cluster 3.x With Sun StorEdge D2 Array Manual

Sun Cluster 3.x With Sun StorEdge MultiPack Enclosure Manual

Sun Cluster 3.x With Sun StorEdge Netra D130 or StorEdge S1 Enclosure Manual

Sun Cluster 3.x With Sun StorEdge T3 or T3+ Array Manual

- The Sun Cluster 3.1 Reference Collection, which contains the following manual:

Sun Cluster 3.1 Reference Manual

- The Sun Cluster 3.1 Data Services Collection: For a listing of the manuals contained in this collection, see the *Sun Cluster 3.1 Data Service 5/03 Release Notes*.

In addition, the `docs.sun.comSM` web site enables you to access Sun Cluster documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at the following Web site:

<http://docs.sun.com>

Documentation Issues

This section discusses known errors or omissions for documentation, online help, or man pages and steps to correct these problems.

Software Installation Guide

This section discusses known errors or omissions from the *Sun Cluster 3.1 Software Installation Guide*.

Quorum-Device Connection

In the *Sun Cluster 3.1 Software Installation Guide*, the following statement about quorum devices is incorrect:

Connection - Do not connect a quorum device to more than two nodes.

The statement should instead read as follows:

Connection – You must connect a quorum device to at least two nodes.

Node Authentication For `scvxinstall` Is Not Required

When you use the `scvxinstall` command to install VERITAS Volume Manager (VxVM), it is no longer necessary to first add the node to the cluster node authentication list. When you perform the procedures in “How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk” or “How to Install VERITAS Volume Manager Software Only”, ignore Step 3, “Add all nodes in the cluster to the cluster node authentication list.”

Upgrade Procedure Refers to Unavailable `scsetup` Functionality

In “How to Prepare the Cluster for Upgrade” in *Sun Cluster 3.1 Software Installation Guide*, the procedure states that, if you are upgrading from Sun Cluster 3.0 5/02 software, you can use the `scsetup` utility to disable resources rather than use the `scswitch` command. This statement is incorrect and should be ignored.

SunPlex Manager Online Help

A note that appears in the Oracle data service installation procedure is incorrect.

Incorrect:

Note: If no entries exist for the `shmsys` and `semsys` variables in the `/etc/system` file when SunPlex Manager packages are installed, default values for these variables are automatically put in the `/etc/system` file. The system must then be rebooted. Check Oracle installation documentation to verify that these values are appropriate for your database.

Correct:

Note: If no entries exist for the `shmsys` and `semsys` variables in the `/etc/system` file when you install the Oracle data service, default values for these variables can be automatically put in the `/etc/system` file. The system must then be rebooted. Check Oracle installation documentation to verify that these values are appropriate for your database.

System Administration Guide

This section discusses errors and omissions from the *Sun Cluster 3.1 System Administration Guide*.

Simple Root Disk Groups With VERITAS Volume Manager

Simple root disk groups are not supported as disk types with VERITAS Volume Manager on Sun Cluster software. As a result, if you perform the procedure “How to Restore a Non-Encapsulated root (/) File System (VERITAS Volume Manager)” in the *Sun Cluster 3.1 System Administration Guide*, you should ignore Step 9, which asks you to determine if the root disk group (`rootdg`) is on a single slice on the root disk. You would complete Step 1 through Step 8, skip Step 9, and proceed with Step 10 to the end of the procedure.

Changing the Number of Node Attachments to a Quorum Device

When increasing or decreasing the number of node attachments to a quorum device, the quorum vote count is not automatically recalculated. You can re-establish the correct quorum vote if you remove all quorum devices and then add them back into the configuration.

Data Services Collection

Errors and omissions related to the Data Services documentation are described in the *Sun Cluster 3.1 Data Service 5/03 Release Notes*.

Man Pages

Sun Cluster 3.0 Data Service Man Pages

To display Sun Cluster 3.0 data service man pages, install the latest patches for the Sun Cluster 3.0 data services that you installed on Sun Cluster 3.1 software. See “Patches and Required Firmware Levels” on page 24 for more information.

After you have applied the patch, access the Sun Cluster 3.0 data service man pages by issuing the `man -M` command with the full man page path as the argument. The following example opens the Apache man page.

```
% man -M /opt/SUNWscapc/man SUNW.apache
```

Consider exporting your `MANPATH` to enable access to Sun Cluster 3.0 data service man pages without specifying the full path. The following example describes command input for adding the Apache man page path to your `MANPATH` and displaying the Apache man page.

```
% MANPATH=/opt/SUNWscapc/man:$MANPATH; export MANPATH
% man SUNW.apache
```

scconf_transp_adap_wrsm(1M)

The following `scconf_transp_adap_wrsm(1M)` man page replaces the existing `scconf_transp_adap_wrsm(1M)` man page.

NAME

`scconf_transp_adap_wrsm.1m`- configure the `wrsm` transport adapter

DESCRIPTION

`wrsm` adapters may be configured as cluster transport adapters. These adapters can only be used with transport types `dlpi`.

The `wrsm` adapter connects to a transport junction or to another `wrsm` adapter on a different node. In either case, the connection is made through a transport cable.

Although you can connect the `wrsm` adapters directly by using a point-to-point configuration, Sun Cluster software requires that you specify a transport junction, a virtual transport junction. For example, if `node1:wrsm1` is connected to `node2:wrsm1` directly through a cable, you must specify the following configuration information.

```
node1:wrsm1 <--cable1--> Transport Junction sw_wrsm1 <--cable2--> node2:wrsm1
```

The transport junction, whether a virtual switch or a hardware switch, must have a specific name. The name must be `sw_wrsmN` where the adapter is `wrsmN`. This requirement reflects a Wildcat restriction that requires that all `wrsm` controllers on the same Wildcat network have the same instance number.

When a transport junction is used and the endpoints of the transport cable are configured using `scconf`, `scinstall`, or other tools, you are asked to specify a port name on the transport junction. You can provide any port name, or accept the default, as long as the name is unique for the transport junction.

The default sets the port name to the node ID that hosts the adapter at the other end of the cable.

Refer to `scconf(1M)` for more configuration details.

There are no user configurable properties for cluster transport adapters of this type.

SEE ALSO

`scconf(1M)`, `scinstall(1M)`, `wrsmconf(1M)`, `wrsmstat(1M)`, `wrsm(7D)`, `wrsm(7D)`

`scconf_transp_adap_sci(1M)`

The `scconf_transp_adap_sci(1M)` man page states that SCI transport adapters can be used with the `rsm` transport type. This support statement is incorrect. SCI transport adapters do *not* support the `rsm` transport type. SCI transport adapters support the `dlpi` transport type only.

`scconf_transp_adap_sci(1M)`

The following sentence clarifies the name of an SCI-PCI adapter. This information is not currently included in the `scconf_transp_adap_sci(1M)` man page.

New Information:

Use the name `sciN` to specify an SCI adapter.

`scgdevs(1M)`

The following paragraph clarifies behavior of the `scgdevs` command. This information is not currently included in the `scgdevs(1M)` man page.

New Information:

`scgdevs(1M)` called from the local node will perform its work on remote nodes asynchronously. Therefore, command completion on the local node does not necessarily mean it has completed its work cluster wide.

`SUNW.sap_ci(5)`

- There is an error in the Name section. The Name section should read as follows:
sap_ci, SUNW.sap_ci and SUNW.sap_ci_v2 - Resource type implementations for Sun Cluster HA for SAP central instance.
- There is an error in the Description section. The Description section should read as follows:
The Resource Group Manager (RGM) manages the SAP data service for Sun Cluster software. Configure the Sun Cluster HA for SAP central instance as a logical-hostname resource and an SAP central instance resource.

SUNW.sap_as(5)

- There is an error in the Name section. The Name section should read as follows:
sap_as, SUNW.sap_as - Resource type implementation for Sun Cluster HA for SAP as a failover data service.
sap_as, SUNW.sap_as_v2 - Resource type implementation for Sun Cluster HA for SAP as a failover data service or a scalable data service.
- There is an error in the Description section. The Description section should read as follows:
The Resource Group Manager (RGM) manages the SAP data service for Sun Cluster software. If you are setting up the Sun Cluster HA for SAP application server as a failover data service configure it as a logical-hostname resource and an SAP application-server resource. If you are setting up the Sun Cluster HA for SAP application-server as a scalable data service configure it as a scalable SAP application-server resource.

rg_properties(5)

The following new resource group property should be added to the `rg_properties(5)` man page.

`Auto_start_on_new_cluster`

This property controls whether the Resource Group Manager starts the resource group automatically when a new cluster is forming.

The default is `TRUE`. If set to `TRUE`, the Resource Group Manager attempts to start the resource group automatically to achieve `Desired primaries` when all nodes of the cluster are simultaneously rebooted. If set to `FALSE`, the Resource Group does not start automatically when the cluster is rebooted. If set to `FALSE`, the Resource Group does not automatically start when the cluster is rebooted. The resource group will remain offline until the first time it is manually switched online using `scswitch (1M)`. After that, it will resume normal failover behavior.

Category: Optional Default: True Tunable: Any time

rt_properties(5)

In this release, the current `API_version` has been incremented to 3 from its previous value of 2. To prevent a resource type from registering on an earlier version of Sun Cluster software, declare `API_version=3`. For more information, see `rt_reg (4)` and `rt_properties (5)`.

Sun Cluster Installation and Configuration Worksheets

This appendix provides worksheets to plan various components of your cluster configuration and examples of completed worksheets for your reference. See “Installation and Configuration Worksheets” in *Sun Cluster 3.1 Data Service 5/03 Release Notes* for configuration worksheets for resources, resource types, and resource groups.

Installation and Configuration Worksheets

If necessary, make additional copies of a worksheet to accommodate all the components in your cluster configuration. Follow planning guidelines in the *Sun Cluster 3.1 Software Installation Guide* to complete these worksheets. Then refer to your completed worksheets during cluster installation and configuration.

Note – The data used in the worksheet examples is intended as a guide only. The examples do not represent a complete configuration of a functional cluster.

The following table lists the planning worksheets and examples provided in this appendix, as well as the titles of sections in “Planning the Sun Cluster Configuration” in *Sun Cluster 3.1 Software Installation Guide* that contain related planning guidelines.

TABLE A-1 Cluster Installation Worksheets and Related Planning Guidelines

Worksheet	Example	Section Titles of Related Planning Guidelines
“Local File System Layout Worksheet” on page 42	“Example: Local File System Layout Worksheets, With and Without Mirrored Root” on page 43	“System Disk Partitions” “Mirroring the Root Disk”
“Cluster and Node Names Worksheet” on page 44	“Example: Cluster and Node Names Worksheet” on page 45	“Cluster Name” “Node Names” “Private Network” “Private Hostnames”
“Cluster Interconnect Worksheet” on page 46	“Example: Cluster Interconnect Worksheet” on page 47	“Cluster Interconnect”
“Public Networks Worksheet” on page 48	“Example: Public Networks Worksheet” on page 49	“Public Networks” “IP Network Multipathing Groups”
“Local Devices Worksheets” on page 50	“Example: Local Devices Worksheets” on page 51	---
“Disk Device Group Configurations Worksheet” on page 52	“Example: Disk Device Group Configurations Worksheet” on page 53	“Disk Device Groups” “Planning Volume Management”

TABLE A-1 Cluster Installation Worksheets and Related Planning Guidelines (Continued)

Worksheet	Example	Section Titles of Related Planning Guidelines
"Volume Manager Configurations Worksheet" on page 54	"Example: Volume Manager Configurations Worksheet" on page 55	"Planning Volume Management" Your volume manager documentation
"Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)" on page 56	"Example: Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)" on page 57	"Planning Volume Management" <i>Solstice DiskSuite 4.2.1 Installation and Product Notes</i> or <i>Solaris Volume Manager Administration Guide</i>

Local File System Layout Worksheet

Node name: _____

TABLE A-2 Local File Systems With Mirrored Root Worksheet

Volume Name	Component	Component	File System	Size
			/	
			swap	
			/globaldevices	

TABLE A-3 Local File Systems with Non-Mirrored Root Worksheet

Device Name	File System	Size
	/	
	swap	
	/globaldevices	

Example: Local File System Layout Worksheets, With and Without Mirrored Root

Node name: **phys-schost-1**

TABLE A-4 Example: Local File Systems With Mirrored Root Worksheet

Volume Name	Component	Component	File System	Size
d1	c0t0d0s0	c1t0d0s0	/	6.75 GB
d2	c0t0d0s1	c1t0d0s1	swap	750 MB
d3	c0t0d0s3	c1t0d0s3	/globaldevices	512 MB
d7	c0t0d0s7	c1t0d0s7	SDS replica	20 MB

TABLE A-5 Example: Local File Systems With Non-Mirrored Root Worksheet

Device Name	File System	Size
c0t0d0s0	/	6.75 GB
c0t0d0s1	swap	750 MB
c0t0d0s3	/globaldevices	512 MB
c0t0d0s7	SDS replica	20 MB

Cluster and Node Names Worksheet

TABLE A-6 Cluster and Node Names Worksheet

Component	Default	Actual
Cluster name		
Private network address	172.16.0.0	_____._____.0.0
Private network mask	255.255.0.0	255.255._____._____
First-Installed node name		
Private hostname	clusternode_____-priv	
Additional node name		
Private hostname	clusternode_____-priv	
Additional node name		
Private hostname	clusternode_____-priv	
Additional node name		
Private hostname	clusternode_____-priv	

Example: Cluster and Node Names Worksheet

TABLE A-7 Example: Cluster and Node Names Worksheet

Component	Default	Actual
Cluster name		sc-cluster
Private network address	172.16.0.0	172.16.0.0
Private network mask	255.255.0.0	255.255.0.0
First-Installed node name		phys-schost-1
Private hostname	clusternode1-priv	phys-schost-1-priv
Additional node name		phys-schost-2
Private hostname	clusternode2-priv	phys-schost-2-priv
Additional node name		
Private hostname	clusternode____-priv	
Additional node name		
Private hostname	clusternode____-priv	

Cluster Interconnect Worksheet

TABLE A-8 Cluster Interconnect Worksheet

Node Name	Adapter Name	Transport Type	Junction Name	Junction Type	Port Name

Example: Cluster Interconnect Worksheet

TABLE A-9 Example: Cluster Interconnect Worksheet

Node Name	Adapter Name	Transport Type	Junction Name	Junction Type	Port Name
phys-schost-1	hme0	dlpi	switch1	switch	1
phys-schost-1	hme1	dlpi	switch2	switch	1
phys-schost-2	hme0	dlpi	switch1	switch	2
phys-schost-2	hme1	dlpi	switch2	switch	2

Public Networks Worksheet

TABLE A-10 Public Networks Worksheet

Component	Name
Node name	
Primary hostname	
IP Network Multipathing group	
Adapter name and test IP address	
Backup adapter and test IP address (optional)	
Network name	
Secondary hostname	
IP Network Multipathing group	
Adapter name and test IP address	
Backup adapter and test IP address (optional)	
Network name	
Secondary hostname	
IP Network Multipathing group	
Adapter name and test IP address	
Backup adapter and test IP address (optional)	
Network name	
Secondary hostname	
IP Network Multipathing group	
Adapter name and test IP address	
Backup adapter and test IP address (optional)	
Network name	

Example: Public Networks Worksheet

TABLE A-11 Example: Public Networks Worksheet

Component	Name
Node name	phys-schost-1
Primary hostname	phys-schost-1
IP Network Multipathing group	ipmp0
Adapter name and test IP address	qfe0, schost-85-t-1a
Backup adapter and test IP address (optional)	qfe4, schost-85-t-1b
Network name	net-85
Secondary hostname	phys-schost-1-86
IP Network Multipathing group	ipmp1
Adapter name and test IP address	qfe1, schost-86-t-1a
Backup adapter and test IP address (optional)	qfe5, schost-86-t-1b
Network name	net-86
Secondary hostname	
IP Network Multipathing group	
Adapter name and test IP address	
Backup adapter and test IP address (optional)	
Network name	
Secondary hostname	
IP Network Multipathing group	
Adapter name and test IP address	
Backup adapter and test IP address (optional)	
Network name	

Local Devices Worksheets

Node name: _____

TABLE A-12 Local Disks Worksheet

Local Disk Name	Size

TABLE A-13 Other Local Devices Worksheet

Device Type	Name

Example: Local Devices Worksheets

Node name: `phys-schost-1`

TABLE A-14 Example: Local Disks Worksheet

Local Disk Name	Size
<code>c0t0d0</code>	2G
<code>c0t1d0</code>	2G
<code>c1t0d0</code>	2G
<code>c1t1d0</code>	2G

TABLE A-15 Example: Other Local Devices Worksheet

Device Type	Name
tape	<code>/dev/rmt/0</code>

Disk Device Group Configurations Worksheet

Volume manager (circle one):

Solstice DiskSuite | Solaris Volume Manager | VxVM

TABLE A-16 Disk Device Groups Worksheet

Disk Group/ Diskset Name	Node Names (indicate priority if ordered list)	Ordered priority? (circle one)	Failback? circle one)
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No

Example: Disk Device Group Configurations Worksheet

Volume manager (circle one):

Solstice DiskSuite

TABLE A-17 Example: Disk Device Groups Configurations Worksheet

Disk Group/ Diskset Name	Node Names (indicate priority if ordered list)	Ordered priority? (circle one)	Failback? circle one)
dg-schost-1	1) phys-schost-1, 2) phys-schost-2	Yes	Yes
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No

Example: Volume Manager Configurations Worksheet

Volume manager (circle one):

Solstice DiskSuite

TABLE A-19 Example: Volume Manager Configurations Worksheet

Name	Type	Component	Component
dg-schost-1/d0	trans	dg-schost-1/d1	dg-schost-1/d4
dg-schost-1/d1	mirror	c0t0d0s4	c4t4d0s4
dg-schost-1/d4	mirror	c0t0d2s5	d4t4d2s5

Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)

TABLE A-20 Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)

File System	Metatrans	Metamirrors		Submirrors		Hot Spare Pool	Physical Device	
		(Data)	(Log)	(Data)	(Log)		(Data)	(Log)

Example: Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)

TABLE A-21 Example: Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)

File System	Metatrans	Metamirrors		Submirrors		Hot Spare Pool	Physical Device	
		(Data)	(Log)	(Data)	(Log)		(Data)	(Log)
/A	d10	d11		d12, d13		hsp000	c1t0d0s0, c2t0d1s0	
			d14		d15	hsp006		c1t0d1s6, c2t1d1s6

